

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problems Mailbox.**

THIS PAGE BLANK (USPTO)

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 690 382 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
03.01.1996 Bulletin 1996/01

(51) Int. Cl.<sup>6</sup>: G06F 12/02

(21) Application number: 95109764.1

(22) Date of filing: 23.06.1995

(84) Designated Contracting States:  
DE FR GB NL SE

(30) Priority: 01.07.1994 US 269585

(71) Applicant: SUN MICROSYSTEMS, INC.  
Mountain View, CA 94043 (US)

(72) Inventors:

- Priem, Curtis  
Fremont, California 94536 (US)
- Nishtala, Satyanarayana  
Cupertino, California 95014 (US)
- Lavelle, Michael G.  
Sartoga, California 95070 (US)

- Webber, Thomas  
Cambridge, Massachusetts 02138 (US)
- Lenoski, Daniel E.  
San Jose, California 95129 (US)
- Mehring, Peter A.  
Sunnyvale, California 94087 (US)
- Moffat, Guy  
Mountain View, California 94043 (US)
- Owen, Christopher R.  
Los Gatos, California 95032 (US)

(74) Representative: Sparing - Röhl - Henseler  
Patentanwälte  
D-40237 Düsseldorf (DE)

## (54) Computer system with a multiplexed address bus and pipelined write operations

(57) A computer system with a multiplexed address bus that is shared by both system memory and by slave devices is described. The slave devices are incorporated into an existing system memory configuration by providing a bus controller to execute a two-cycle address sequence on the multiplexed address bus. The address sequence is followed by a transfer of data. A random latency can exist between the time of receiving address information and the time of receiving data corresponding to the address information. This random latency can be exploited by the system CPU for other computational purposes. The bus controller of the system executes multiple, or pipelined, data writes to the bus before an acknowledgement for the first data write is received. In this scheme, the acknowledgement for the first data write is typically sent during the same time period that the subsequent data writes are being received. Consequently, data transfer acknowledgements overlap data writes. This overlapping operation allows the bus to be completely utilized during write operations, thereby improving data bandwidth.




FIG. 2

EP 0 690 382 A2

## Description

### Brief Description of the Invention

5 This invention relates generally to the transfer of data in computer systems. More particularly, this invention relates to a bus used between a central processing unit, system memory, and one or more slave devices, which uses multiplexed address cycles, and which executes pipelined write operations to reduce data transfer acknowledgement overhead.

### Background of the Invention

10 A computer system relies upon a high-speed memory bus to transfer data between system memory and a central processing unit (CPU). A separate lower speed input/output (I/O) bus is typically used for I/O operations between a CPU and a slave device, such as a print driver. Prior art computer systems have used lower speed I/O buses for data transfer between system memory and a CPU. This technique considerably increases the time required for reading and writing data to memory. Thus, prior art computer systems have provided a high-speed memory bus for the system memory and a separate low-speed I/O bus for slave devices. It would be highly desirable to eliminate this redundancy in hardware and utilize a single high-speed memory bus for data transfers between system memory and a CPU, and for controlling operations between a CPU and a slave device. Such a system should not require a reconfiguration of system memory. In other words, the slave device should operate in conjunction with the existing system memory access architecture.

20 Prior art techniques for transferring data between a CPU and a slave device on a low speed I/O bus is a two step process. First, the CPU, also called a master device, generates a write request, a memory mapped I/O address, and data onto the I/O bus. A slave device, such as a print driver, reads the information, and then sends an acknowledgement back to the master device. The time between the write request and the data acknowledgement is the data transfer acknowledgement latency.

25 The data bandwidth of prior art computer buses is limited by this data transfer acknowledgement latency. The data transfer acknowledgement latency problem is more fully appreciated with reference to Figure 1. A write request is sent from a master device. An acknowledgement signal signifying the receipt of the data is then sent from a slave device several cycles later. Only after this acknowledgement signal is received can the master device issue another write request, as shown in Figure 1. On a synchronous bus, this data transfer acknowledgement latency doubles or triples the number of clock cycles required to send a single write request.

30 To mitigate the bandwidth expense associated with data transfer acknowledgement latencies, burst mode data transfers were developed. In a burst mode data transfer, only one acknowledgement is required for a large block of data. Thus, burst mode data transfer amortizes the data transfer acknowledgement latency over larger blocks of write data.

35 Unfortunately, the burst mode data transfer technique has associated problems. First, the data transfer acknowledgement latency is just as long as in the traditional data transfer case shown in Figure 1. Another problem is that the burst mode data transfer requires that the same amount of data always be transferred. Frequently, the amount of data to be sent is less than the required amount for burst mode data transfer. This problem may be alleviated by special techniques for masking the unused portion of the burst mode data. However, this results in added system complexity and expense.

40 Another problem with burst mode data transfers is that they require a fixed byte order. It is frequently desirable to transfer data without conforming to a fixed byte order, for instance in an overlapping copy operation.

45 Burst mode data transfer techniques have a restricted address range. That is, only addresses within the burst size can be accessed in the burst mode. It is desirable to access any address in any order. This capability is especially important in graphics applications. For example, when drawing vectors and triangles, where addresses are localized, but they are in random directions.

Thus, it would be highly desirable to provide a computer system data bus that reduces data transfer acknowledgement latency overhead. Such a system should avoid the problems attendant with burst mode data transfers.

### Summary of the Invention

50 The invention according to one embodiment is a computer system with a multiplexed address bus that is shared by both system memory and by slave devices. The slave devices are incorporated into an existing system memory configuration by providing a bus controller to execute a two-cycle address sequence on the multiplexed address bus. Each address sequence includes a first write cycle in which high-order address bits are launched on the bus and latched by a selected slave device. The address sequence also includes a second address cycle in which the remaining low-order address bits and byte mask bits are sent to the selected slave device. In subsequent address sequences, the high-order address bits are only resent if their content changes. Consequently, most address sequences from the bus controller to a slave device include only a single address cycle.

The address sequence is followed by a transfer of data. Thus, the system has a slave device which utilizes decoupled address and data. A random latency can exist between the time of receiving address information and the time of receiving data corresponding to the address information. This random latency can be exploited by the system CPU for other computational purposes.

5 The bus controller of the system executes multiple, or pipelined, data writes to the bus before an acknowledgement for the first data write is received. In this scheme, the acknowledgement by the slave device for the first data write is typically sent during the same time period that subsequent data writes are being executed by the master device. Consequently, data transfer acknowledgements overlap data writes. This overlapping operation allows the bus to be completely utilized during write operations, thereby improving data bandwidth.

#### 10 Brief Description of the Drawings

For a better understanding of the nature and objects of the invention, reference should be made to the following detailed description taken in conjunction with the accompanying drawings, in which:

15 FIGURE 1 is a set of timing signals used to illustrate the prior art problem of data transfer acknowledgment latency.  
FIGURE 2 is a diagram of the computer system in accordance with the invention.  
FIGURE 3 is a timing diagram illustrating the multiplexed address and the pipelined write operations of the invention.  
FIGURE 4 depicts a slave device that may be used in accordance with the invention.  
20 FIGURE 5 is a timing diagram illustrating the split read operations of the invention.

Like reference numerals refer to corresponding parts throughout the several views of the drawings.

#### Detailed Description of the Invention

25 The present invention is directed toward a computer system with a multiplexed address bus that is shared by both system memory and by slave devices. The single bus eliminates hardware redundancy. In addition, it allows the slave devices to access the large data fields, 64 bits in this example, used by the system memory. In accordance with the invention, slave devices are incorporated into an existing system memory configuration by providing a two-cycle address sequence. Each address sequence includes a first write cycle in which high-order address bits are launched on the bus and latched by a slave device. The address sequence also includes a second address cycle in which the remaining low-order address bits and byte mask bits are sent to the slave device. In subsequent address sequences, the high-order address bits are only re-sent if their content changes. Consequently, most address sequences to a slave device only require a single address cycle conveying the low-order address bits. This feature improves the data bandwidth to the slave device. This is similar to fast page mode in memory systems, but it is believed to be novel when used in relation to a slave device.

30 The address sequence is followed sometime later by a transfer of data. The system of the present invention can have a random latency between the time of receiving address information and the time of receiving data corresponding to the address information. This random latency can be exploited for other computational purposes. The ability to provide a random latency between address and data write operations is believed to be novel.

40 The present invention also improves the data bandwidth on data buses of a computer system by allowing for multiple, or pipelined, data writes to the data bus before an acknowledgement for the first data write is received. In this scheme, the acknowledgement for the first data write is typically sent during the same time period that subsequent data writes are being executed. Consequently, data transfer acknowledgements overlap data writes. This overlapping operation allows the data bus to be completely utilized during write operations, thereby improving data bandwidth.

45 The salient features of the invention have now been described. At this time, attention turns to a more detailed consideration of the invention. First, the architecture of the multiplexed bus of the invention will be described. Thereafter, attention will turn to the pipelined writing operation of the invention.

50 Figure 2 illustrates a computer system 20 incorporating the multiplexed bus 22 of the invention. The multiplexed bus 22 includes control lines 22A and address/data lines 22B. The multiplexed bus 22 is used to incorporate one or more slave devices 24 into the system 20. The slave devices 24 are controlled by a bus controller 26. For the sake of simplicity, the invention will be described in reference to only one slave device 24. However, it should be appreciated that multiple slave devices may be used.

55 The control lines 22A, bus controller 26, and slave device 24 are incorporated into an existing system that includes a central processing unit (CPU) 28, a memory controller 30, memory control lines 32, and system memory 34. The address/data lines 22B constitute the existing address/data lines used by system memory 36. Thus, it can be appreciated that the multiplexed bus of the present invention exploits existing system memory 36 resources and adds additional control lines so that a slave device can be used. In this way, the slave device 24 does not require separate address and data buses, thereby reducing system hardware requirements. This feature is also advantageous because the absence

of a second bus reduces the number of pins that the CPU must support, an important packaging consideration. Moreover, the slave device can access the large bandwidth data of the address/data lines 22B. Consequently, a high data throughput device, such as a graphics card, may be used as the slave device 24.

The invention is directed toward a multiplexed bus 22 that exploits existing system resources. The existing system resource including the CPU 28, memory controller 30, and system memory 34 are known in the art and need not be discussed in detail. In addition, the contention operation between the bus controller 26 and the memory controller 30 for the address/data lines 22B may be operated in accordance with any number of contention resolution techniques. It is only important to recognize that the address/data lines 22B are used by the system memory 34 and the slave device 24 at different times.

The address/data lines 22B are in the configuration of an existing computer bus. Figure 2 indicates that the address/data lines include the following lines: AB[10:0], WRITE\_L, and DB[63:0]. The AB[10:0] symbol indicates that there are 11 address lines, the DB[63:0] symbol indicates that there are 64 data lines, and the WRITE\_L line is used to distinguish between a read or a write operation.

The multiplexed bus of the invention uses these existing lines and additionally provides for a set of control lines 22A so that the slave device 24 may be serviced. Note in Figure 2 that the control lines 22A include the following lines: AEN (address enable), AB[14:11], S\_REPLY[1:0], P\_REPLY[1:0], LO\_ADDR (low address), INT\_L, CLK, RESET\_L, and DEV\_PRES\_L. The control lines 22A also include test signal lines and power and ground lines, which are omitted for the sake of simplicity.

The operation of each line of the multiplexed bus may be characterized as follows. AEN is an address enable signal. When AEN is asserted (high), it indicates that there is a command request and there is valid data on the AB bus and the LO\_ADDR and WRITE\_L lines. The AEN signal also qualifies read cycles. The address/byte mask bus, AB[14:0], contains the multiplexed physical address indicating where the data should be written to or read from. In addition, it contains a byte mask. This address/byte mask data is multiplexed over two address cycles, as will be described below.

The system reply signals, S\_REPLY[1:0], indicate that the bus has been selected and the type of access. The port reply signals, P\_REPLY[1:0], indicate that the data has been processed (removed) from the write buffer on writes or that the read data is available in the read latch.

The low address signal, LO\_ADDR, is qualified by the AEN signal. The signal defines which address cycle the controller is sending on the address bus and write lines. When driven to 0, this signal indicates address cycle 0 (high address bits). When driven to 1, this signal indicates address cycle 1 (low address bits and byte mask).

The bus interrupt signal, INT\_L, is used by a slave 24 to signal the controller 26 that some event has occurred and requires servicing. Before asserting the interrupt, the slave must set a bit in an internal register to indicate the cause of the interrupt. The interrupt is serviced by the CPU 28 reading this bit. The bit may be cleared by reading it or by an explicit write by the CPU 28. Once cleared, the slave 24 must then unassert the interrupt signal.

The signal, DEV\_PRES\_L, is used to identify the presence of a device. Specifically, the pin is used to signify that a slave device 24 is present at a physical location.

The bus reset signal, RESET\_L, properly initializes all bus slaves after power-up and system reset. Upon detecting the assertion of RESET\_L, a bus slave 24 must perform whatever internal operations are required to initialize itself.

The WRITE\_L signal is qualified by AEN and is only valid during address cycle 1. When driven to 0, this signal indicates a write request. When driven to 1, this signal indicates a read request.

The data lines, DB[63:0], are used to transfer data between the CPU 28 and the bus slave 24. The data lines support four primary data formats: (1) bytes, which consist of eight data bits, (2) half words, which consist of 16 data bits, (3) words, which consist of 32 data bits, and (4) double words, which consist of 64 data bits. By convention, the least-significant bit of the data bus is DB[0], while the most-significant bit is DB[63].

The operation of each of the described lines of the multiplexed bus 22 is summarized in Table I.

TABLE I

| SIGNAL NAME | I/O | DESCRIPTION        | DRIVEN BY        |
|-------------|-----|--------------------|------------------|
| AEN         | I   | Address Enable     | Controller       |
| AB          | I   | Address/Byte Mask  | Controller       |
| S_REPLY     | I   | System Reply       | Controller       |
| P_REPLY     | O   | Port Reply         | Slave            |
| LO_ADDR     | I   | Low Address Select | Controller       |
| INT_L       | O   | Interrupt          | Slave            |
| CLK         | I   | Clock              | Controller       |
| DEV_PRES_L  | O   | Device Present     | Slave            |
| RESET_L     | I   | Reset              | Controller       |
| WRITE_L     | I   | Read/Write Select  | Controller       |
| DB[63:0]    | I/O | Data lines         | Controller/Slave |

The multiplexed bus 22 of the invention executes a two-cycle address sequence operation that is described in reference to the foregoing signals and the following table.

TABLE II

| ADDRESS BUS  | CYCLE 0 | CYCLE 1 |
|--|---------|---------|
| LO_ADDR  | 0       | 1       |
| AB[14]   | PA[27]  | BM[3]   |
| AB[13]   | PA[26]  | BM[2]   |
| AB[12]   | PA[25]  | BM[1]   |
| AB[11]   | PA[24]  | BM[0]   |
| AB[10]   | PA[23]  | PA[13]  |
| AB[9]  | PA[22]  | PA[12]  |
| AB[8]  | PA[21]  | PA[11]  |
| AB[7]  | PA[20]  | PA[10]  |
| AB[6]  | PA[19]  | PA[9]   |
| AB[5]  | PA[18]  | PA[8]   |
| AB[4]  | PA[17]  | PA[7]   |
| AB[3]  | PA[16]  | PA[6]   |
| AB[2]  | PA[15]  | PA[5]   |
| AB[1]  | PA[14]  | PA[4]   |
| AB[0]  | X       | PA[3]   |
| WRITE_L  | X       | R/W     |
| PA = Physical Address<br>BM[3:0] = Byte Mask Bits<br>R/W = Read or Write<br>X = Don't Care |         |         |

A line appears in Table II to demark that the first set of signals are applied to the control lines 22A, and the second set of signals are applied to the address/data lines 22B. During cycle 0, the slave device 24 latches the high-order address bits (PA[27]-PA[14]). These latched values are then combined with the low-order address bits (PA[13]-PA[3]) in each subsequent address cycle. Another address cycle 0 is only needed when the data access crosses the "page boundary" and the high-order address bits need to be modified.

Note in cycle 1 that byte mask values are latched by the slave device 24. These byte mask values are used to identify where the data is in the 64 bit field. As indicated above, the data lines supports four primary data formats: (1) bytes, which consist of eight data bits, (2) half words, which consist of 16 data bits, (3) words, which consist of 32 data bits, and (4) double words, which consist of 64 data bits. Unlike a prior art burst mode data transfer, the present invention allows any sequence of addresses within the low address range. This is helpful in overlapping copy operations where the data may have to be read and written in reverse order. It also helps in vector and triangle graphics applications where access is localized but in random directions.

The two-cycle address sequence operation of the multiplexed bus 22 is more fully appreciated with reference to Figure 3. Figure 3 shows that the address enable signal AEN initiates the writing of the high-order bits to AB[14:0] in a first cycle. Thereafter, in response to the low address signal LO\_ADDR, the low-order bits are written. Note that since the high-order bits do not change, they are not re-sent and a sequence of low order bits are written.

Figure 4 shows the processing by the slave 24 in response to the two-cycle address sequence operation of the multiplexed bus 22. In the initial cycle of the write sequence, the high-order address bits are sent. Specifically, the high-order address signals (AB[14:11]) of the high-order address bits are latched into latch 40 over the control lines 22A, and the low-order address bits (AB[10:0]) of the high-order address signals are latched into latch 40 over the address/data lines 22B.

In the next cycle of the address sequence, byte masks are loaded by the control lines 22A into the latch 40 by lines AB[14:11], and the remaining low-order address bits (AB[10:0]) of the high-order address signals are loaded by the address/data lines 22B into latch 40. All of the address data has now been obtained. When subsequent address writes are initiated, the data in the latch 40 is sent to a standard data storage structure 43, such as a first-in-first-out (FIFO) data storage structure. Figure 4 shows an address FIFO 43 for storing a sequence of addresses.

After each address is written, the data corresponding to the address is received in a data FIFO 42. The control logic 44 coordinates the received address in latch 40 and the received data in FIFO 42. Recall that the address information includes byte mask data that is used by the control logic 44 to select the proper data. Using memory mapped I/O 46, the control logic 44 executes a selected operation in relation to the received data.

The multiplexed bus 22 of the invention has now been fully described. Attention presently turns to the pipelined write operations of the invention.

As previously indicated, the slave device 24 has an address FIFO 43 and a data FIFO 42 that hold transmitted address and data information. Write requests sent to a slave device 24 are saved in the address FIFO 43 and the data FIFO 42. Subsequently, the control logic reads the data from the respective FIFOs 42, 43 and sends an acknowledge signal to the bus controller 26. Thus, in accordance with the invention, a predetermined number of write requests may be performed and loaded into the address FIFO 43 and the data FIFO 42 of the slave device 24.

This operation is shown in relation to Figure 3. In Figure 3 four write requests (A, B, C, D) are launched on the (S\_REPLY) lines. As previously indicated, the S\_REPLY signal goes directly to the slave device 24 over the control lines 22A. Datapath writes are then initiated on the data lines 22B (DB[63:0]) within one cycle of each write request. The acknowledgements on the P\_Reply line, of the control lines 22A, overlap with the write operation. For instance, write request C occurs at the same time as the acknowledgement for write data A.

Thus, it can be appreciated from Figure 3 that sequential write operations are performed before and during the time that acknowledgements are received.

This approach is in contrast to the prior art wherein an acknowledgement is required before each write operation. Thus, in accordance with the invention, a high bandwidth of pipelined write operations is achieved. Note that each write operation has to be acknowledged. In a burst mode data transfer technique a number of writes are sent, but only one acknowledgement is used for the group of writes. The pipelined write operation of the invention does not suffer many of the shortcomings associated with a burst mode data transfer. Specifically, the pipelined write operation does not have to transmit a fixed amount of data, it does not require a fixed byte order, and it does not have a restricted address range. The previously described mask bytes, used in accordance with the invention, provides greater flexibility in data format than is available in a burst mode data transfer operation.

As previously indicated, the invention is implemented by providing each slave device 24 with FIFOs to store the received address and write data. The FIFOs corresponding to the operation in Figure 3 must be capable of holding four writes. The invention can be implemented with a data storage structure of any number of write operations.

The bus controller 26 operates in a standard manner, but also includes a counter to keep track of the number of write operations. When the counter reaches a predetermined write limit, it activates a control circuit that prevents the bus controller 26 from generating writes beyond the capacity of the FIFOs 42, 43. Thus, in the example of Figure 3, after four write operations, the bus controller 26 stops executing data writes to the datapath bus 22B.



In Figure 3, E and F are read cycles. Note that there is a space before read cycle E. It is desirable to design the system 20 so that a read is launched only if there are 1 or zero unacknowledged writes in the FIFOs of the slave 24. Thus, the E write request is launched after the acknowledgement for C is sent on the P\_REPLY line. The system 20 may be designed such that a read can be launched despite a large number of unacknowledged writes in the FIFO. However, if the number is too high, then there is a danger that a read operation will stall the system 20 for an unreasonable period of time.

The read operation is performed when the AEN signal is activated, the "Read Single" code (11) is placed on the S\_REPLY lines, and the low address write on AB[14:0] of cycle 1 are all performed in the same clock cycle, as shown in Figure 3. When the slave 24 is ready to acknowledge the read, the slave 24 simultaneously drives the data on the DB lines, through register 48 shown in Figure 4, and a "Read Single" code (11) on the P\_Reply lines. This mode has the advantage of doing a best-case read in two clock cycles.

To fully exploit the architecture of the present invention, the bus controller 26 and the slave device 24 preferably support an additional type of read operation, referred to herein as a split read. The split read allows the CPU 28 to perform functions, say memory writes, while its waiting for the read data.

The split read feature exploits the fact that in the present invention the address signals are decoupled from the data signals. Thus, with the present invention, the address and data may be sent at different times. However, the address must be sent first. Data can then come any number of clock cycles later.

This operation is more fully appreciated with reference to Figure 5. A split read cycle is identified by the slave 24 when the S\_REPLY signals are idle during the launching address cycle 1. Once the slave 24 has the data for the access, the slave 24 drives a "Read Single" code (11) on the P\_Reply lines. The controller 26 responds with a "Read Single" code (11) on the S\_REPLY lines to tell the slave 24 that the data bus is available. The slave 24 drives the data on the DB lines two clock cycles later, and the transfer is complete.

As shown in Figure 5, the split read takes a minimum of five clock (CLK) cycles. In cycle 0, if required, the bus controller 26 places the address on AB[14:0] and asserts AEN. On the second address cycle (address cycle 1), the bus controller 26 asserts LO\_ADDR. When it is ready for the read (identified by WRITE\_L being unasserted), the slave 24 asserts the P\_REPLY read single (11) code. When it is ready to accept the read data, the bus controller 26 asserts S\_REPLY. Two clock cycles later, the slave 24 places the requested read data on DB[63:0].

The foregoing descriptions of specific embodiments of the present invention are presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, obviously many modifications and variations are possible in view of the above teachings. For example, it should be borne in mind that the invention is applicable to multiple slave devices. In addition, the bus controller 26 need not be embedded within a CPU 28. The embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following Claims and their equivalents.

## Claims

1. A computer system, comprising:
  - a multiplexed address bus;
  - a bus controller to generate an address transfer sequence including a first cycle with high order address bits that are launched on said multiplexed bus and a second cycle with low order address bits that are launched on said multiplexed address bus, and
  - a data transfer sequence executed after a latency period following said address transfer sequence, said data transfer sequence including the launch of a plurality of data bits on said multiplexed address; and
  - a slave device connected to said multiplexed address bus to receive and process said high order address bits, said low order address bits, and said data bits.
2. The computer system of claim 1 further comprising a memory device coupled to said multiplexed address bus.
3. The computer system of claim 2 wherein said multiplexed address bus includes address/data lines, said memory device and said slave device being connected to said address/data lines.
4. The computer system of claim 3 wherein said address/data lines include address lines and data lines.
5. The computer system of one of the claims 1 to 4 wherein said multiplexed address bus includes control lines, said slave device being connected to said control lines.

6. The computer system of claim 5 wherein said control lines include address lines.

7. In a computer system that includes a master device to launch a write request onto a data bus, and a controlled device to receive the write request and launch a write acknowledgement onto the data bus before a subsequent write request can be performed by the master device, the improvement comprising:

a bus controller to launch a sequence of write requests before receiving a write acknowledgement; and  
a slave device to store said sequence of write requests and sequentially provide write acknowledgements on said data bus for each write request of said sequence of write requests.

8. The computer system of claim 7 wherein said slave device includes a data storage structure to store said sequence of write requests, said data storage structure having a corresponding write request limit defining the number of write requests that can be stored.


9. The computer system of claim 7 or 8 wherein said bus controller includes a counter to count each write request and a control circuit responsive to said counter to prevent the number of write requests from exceeding said write request limit of said data storage structure.

10. A method of operating a computer system that includes a master device to launch a write request onto a data bus, and a controlled device to receive the write request and launch a write acknowledgement onto the data bus before a subsequent write request can be performed by the master device, the method comprising the steps of:


launching a sequence of write requests before receiving a write acknowledgement; and  
storing said sequence of write requests and sequentially providing write acknowledgements on said data bus for each write request of said sequence of write requests.

11. The method of claim 10 wherein said storing step includes the step of storing said sequence of write requests up to a write request limit.

12. The method of claim 10 or 11 wherein said launching step includes the steps of  
counting each write request; and  
halting write requests when said write request limit is reached.



**FIG. 1**  
(PRIOR ART)



**FIG. 2**




FIG. 3




FIG. 4




FIG. 5

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 690 382 A3

(12)

## EUROPEAN PATENT APPLICATION

(88) Date of publication A3:  
02.10.1996 Bulletin 1996/40

(51) Int. Cl.<sup>6</sup>: G06F 12/02, G06F 13/16

(43) Date of publication A2:  
03.01.1996 Bulletin 1996/01

(21) Application number: 95109764.1

(22) Date of filing: 23.06.1995

(84) Designated Contracting States:  
DE FR GB NL SE

(30) Priority: 01.07.1994 US 269585

(71) Applicant: SUN MICROSYSTEMS, INC.  
Mountain View, CA 94043 (US)

(72) Inventors:


- Priem, Curtis  
Fremont, California 94536 (US)
- Nishtala, Satyanarayana  
Cupertino, California 95014 (US)
- Lavelle, Michael G.  
Sartoga, California 95070 (US)

- Webber, Thomas  
Cambridge, Massachusetts 02138 (US)
- Lenoski, Daniel E.  
San Jose, California 95129 (US)
- Mehring, Peter A.  
Sunnyvale, California 94087 (US)
- Moffat, Guy  
Mountain View, California 94043 (US)
- Owen, Christopher R.  
Los Gatos, California 95032 (US)

(74) Representative: Sparing - Röhl - Henseler  
Patentanwälte  
Rethelstrasse 123  
40237 Düsseldorf (DE)

(54) **Computer system with a multiplexed address bus and pipelined write operations**

(57) A computer system with a multiplexed address bus that is shared by both system memory and by slave devices is described. The slave devices are incorporated into an existing system memory configuration by providing a bus controller to execute a two-cycle address sequence on the multiplexed address bus. The address sequence is followed by a transfer of data. A random latency can exist between the time of receiving address information and the time of receiving data corresponding to the address information. This random latency can be exploited by the system CPU for other computational purposes. The bus controller of the system executes multiple, or pipelined, data writes to the bus before an acknowledgement for the first data write is received. In this scheme, the acknowledgement for the first data write is typically sent during the same time period that the subsequent data writes are being received. Consequently, data transfer acknowledgements overlap data writes. This overlapping operation allows the bus to be completely utilized during write operations; thereby improving data bandwidth.



EP 0 690 382 A3



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 95 10 9764

| DOCUMENTS CONSIDERED TO BE RELEVANT  |  |  |  |
|--|--|--|--|
| Category   | Citation of document with indication, where appropriate, of relevant passages                            | Relevant to claim                                | CLASSIFICATION OF THE APPLICATION (Int.Cl.6) |
| X  | EP-A-0 020 908 (IBM)<br>* page 3, line 19 - page 5, line 20;<br>figure 1 *                               | 1-6  | G06F12/02<br>G06F13/16                       |
| X  | US-A-4 674 032 (UNISYS)<br>* column 13, line 10 - line 62 *<br>* column 16, line 7 - line 57; figure 9 * | 7-12   |  |
| A  | EP-A-0 288 649 (IBM)<br>* column 3, line 9 - line 33; figure 1 *   | 7-12   |  |
|  |  |  | TECHNICAL FIELDS SEARCHED (Int.Cl.6)         |
|  |  |  | G06F   |
| The present search report has been drawn up for all claims   |  |  |  |
| Place of search<br>THE HAGUE   |  | Date of completion of the search<br>29 July 1996 | Examiner<br>Gill, S                          |
| <p><b>CATEGORY OF CITED DOCUMENTS</b></p> <p>X : particularly relevant if taken alone<br/>Y : particularly relevant if combined with another document of the same category<br/>A : technological background<br/>O : non-written disclosure<br/>P : intermediate document</p> <p>T : theory or principle underlying the invention<br/>E : earlier patent document, but published on, or after the filing date<br/>D : document cited in the application<br/>L : document cited for other reasons<br/>@ : member of the same patent family, corresponding document</p> |  |  |  |

EPO FORM 1503 03.02 (P04C01)





European Patent  
Office

EP 95 10 9764

### CLAIMS INCURRING FEES

The present European patent application comprised at the time of filing more than ten claims.

- ☐ All claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for all claims.
- ☐ Only part of the claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid.
- namely claims:
- ☐ No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

### X LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirement of unity of invention and relates to several inventions or groups of inventions.

namely:

See Sheet 3.

- ☒ All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.
- ☐ Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid.
- namely claims:
- ☐ None of the further search fees has been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims.
- namely claims:



European Patent  
Office

EP 95 10 9764 -B-

|   |                            |
|---|----------------------------|
| X | LACK OF UNITY OF INVENTION |
|---|----------------------------|

The Search Division considers that the present European patent application does not comply with the requirement of unity of invention and relates to several inventions or groups of inventions, namely:

1. Claims 1-6 : Computer system with multiplexed address bus shared by system memory and slave devices.
2. Claims 7-12: Pipelining of data writes on a system bus.